This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
Please do not report the images to the
Image Problem Mailbox.

US005867663A

United States Patent [19]

McClure et al.

[11] Patent Number: 5,867,663

[45] Date of Patent: Feb. 2, 1999

[54] METHOD AND SYSTEM FOR CONTROLLING NETWORK SERVICE PARAMETERS IN A CELL BASED COMMUNICATIONS NETWORK

[75] Inventors: **Robert B. McClure**, Hollis; **Stephen A. Caldara**, Sudbury; **Stephen A. Hauser**, Burlington; **Thomas A. Manning**, Northboro, all of Mass.

[73] Assignees: **Fujitsu Network Communications, Inc.**, Richardson, Tex.; **Fujitsu Limited**, Kawasaki, Japan

[21] Appl. No.: **685,183**

[22] Filed: **Jul. 18, 1996**

Related U.S. Application Data

[60] Provisional application No. 60/001,498 filed Jul. 19, 1995.

[51] **Int. Cl.** [6] **H01J 13/00**

[52] **U.S. Cl.** **395/200.64**; 370/413

[58] **Field of Search** 370/412–413, 370/414–418, 905; 395/200.64

[56] **References Cited**

U.S. PATENT DOCUMENTS

3,804,991	4/1974	Hammond et al.	179/26
3,974,343	8/1976	Cheney et al.	179/18 ES
4,069,399	1/1978	Barrett et al.	179/15 AL
4,084,228	4/1978	Dufond et al.	364/200
4,240,143	12/1980	Bessemer et al.	364/200
4,603,382	7/1986	Cole et al.	364/200
4,715,030	12/1987	Koch et al.	370/85
4,727,537	2/1988	Nichols	370/85
4,737,953	4/1988	Koch et al.	370/94
4,748,658	5/1988	Gopal et al.	379/221
4,797,881	1/1989	Ben-Artzi	370/88
4,821,034	4/1989	Anderson et al.	340/825.8
4,837,761	6/1989	Isono et al.	370/60
4,849,968	7/1989	Turner	370/94
4,870,641	9/1989	Pattavina	370/60
4,872,157	10/1989	Hemmady et al.	370/60

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

484943 3/1992 Japan .

OTHER PUBLICATIONS

Choudhury, A., et al., "Dynamic Queue Length Thresholds in a Shared Memory ATM Switch," Infocom, vol. 2, pp. 679–687, 1996.

Kim, H., "Design and Performance of Multinet Switch: A Multistage ATM Switch Architecture with Partially Shared Buffers," IEEE Transactions on Networking, vol. 2, No. 6, pp. 571–580, Dec. 1994.

(List continued on next page.)

Primary Examiner—Robert B. Harrell
Assistant Examiner—David M. Ovedovitz
Attorney, Agent, or Firm—Weingarten, Schurgin, Gagnebin & Hayes LLP

[57] **ABSTRACT**

A method and apparatus for controlling network service parameters in a cell based communications network. A plurality of input signals are received on input ports of a communications device and a plurality of output signals are sourced from output ports, wherein at least one input signal and at least one output signal comprise, respectively, at least one virtual connection. The communications device is part of a cell based communications network and each virtual connection comprises a series of data cells comprising a header portion and a data portion. A first virtual connection is assigned to an input queue of an input port and an output queue of an output port. Each queue is associated with a first buffer space dedicated to the respective queue. Each queue is also associated with a second buffer space that is shared between the respective queue and other queues. The second buffer space is utilized after the first buffer space becomes full. Data associated with the first virtual connection is queued in the input and output queues. The communications device also includes a switching fabric intermediate the input ports and the output ports. Bandwidth into the switching fabric is allocated by the input port, and bandwidth out of the device is allocated by the output port.

40 Claims, 2 Drawing Sheets

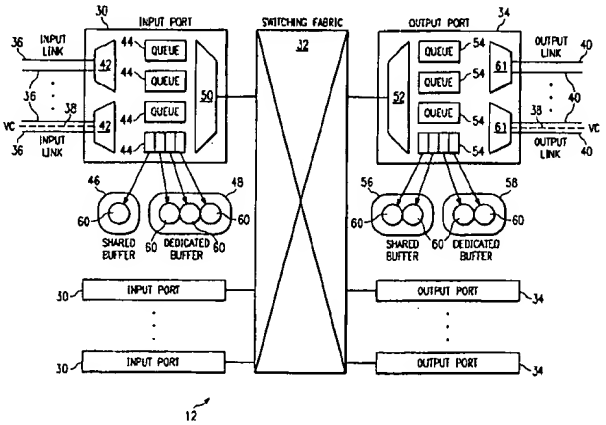

U.S. PATENT DOCUMENTS

4,872,159	10/1989	Hemmady et al.	370/60
4,872,160	10/1989	Hemmady et al.	370/60
4,872,197	10/1989	Pemmaraju	379/93
4,878,216	10/1989	Yunoki	370/60
4,893,302	1/1990	Hemmady et al.	370/60
4,893,307	1/1990	McKay et al.	370/94.1
4,894,824	1/1990	Hemmady et al.	370/58.3
4,897,833	1/1990	Kent et al.	370/85.2
4,897,841	1/1990	Gang, Jr.	370/85.13
4,899,333	2/1990	Roediger	370/60
4,920,531	4/1990	Isono et al.	370/60
4,922,503	5/1990	Leone	370/85.13
4,933,938	6/1990	Sheehy	370/85.13
4,942,574	7/1990	Zelle	370/85.15
4,947,390	8/1990	Sheehy	370/85.13
4,953,157	8/1990	Franklin et al.	370/60
4,956,839	9/1990	Torii et al.	370/60
4,958,341	9/1990	Hemmady et al.	370/60.1
4,979,100	12/1990	Makris et al.	364/200
4,993,018	2/1991	Hajikano et al.	370/60
5,014,192	5/1991	Mansfield et al.	364/200
5,021,949	6/1991	Morten et al.	364/200
5,029,164	7/1991	Goldstein et al.	370/95.1
5,060,228	10/1991	Tsutsui et al.	370/85.13
5,067,123	11/1991	Hyodo et al.	370/58.1
5,070,498	12/1991	Kakuma et al.	370/60
5,083,269	1/1992	Syobatake et al.	395/425
5,084,867	1/1992	Tachibana et al.	370/60
5,084,871	1/1992	Carn et al.	370/94.1
5,090,011	2/1992	Fukuta et al.	370/60
5,090,024	2/1992	Vander Mey et al.	375/1
5,093,827	3/1992	Franklin et al.	370/60.1
5,093,912	3/1992	Dong et al.	395/650
5,115,429	5/1992	Hluchyj et al.	370/84
5,119,369	6/1992	Tanabe et al.	370/60
5,119,372	6/1992	Verbeek	370/85.3
5,128,932	7/1992	Li	370/60
5,130,975	7/1992	Akata	370/60
5,130,982	7/1992	Ash et al.	370/85.7
5,132,966	7/1992	Hayano et al.	370/79
5,146,474	9/1992	Nagler et al.	375/10
5,146,560	9/1992	Goldberg et al.	395/200
5,150,358	9/1992	Punj et al.	370/84
5,151,897	9/1992	Suzuki	370/85.13
5,157,657	10/1992	Potter et al.	370/85
5,163,045	11/1992	Caram et al.	370/60.1
5,163,046	11/1992	Hahne et al.	370/79
5,166,930	11/1992	Braff et al.	370/235
5,179,556	1/1993	Turner	370/94.1
5,179,558	1/1993	Thacker et al.	370/94.3
5,185,743	2/1993	Murayama et al.	370/110.1
5,191,582	3/1993	Upp	370/94.1
5,191,652	3/1993	Dias et al.	395/200
5,193,151	3/1993	Jain	395/200
5,197,067	3/1993	Fujimoto et al.	370/94.1
5,198,808	3/1993	Kudo	340/825.8
5,199,027	3/1993	Barri	370/60
5,239,539	8/1993	Uchida et al.	370/58.3
5,253,247	10/1993	Hirose et al.	370/14
5,253,248	10/1993	Dravida et al.	370/16
5,255,264	10/1993	Cotton et al.	370/24
5,255,266	10/1993	Watanabe et al.	370/60.1
5,257,311	10/1993	Naito et al.	380/48
5,258,979	11/1993	Oomuro et al.	370/95.1
5,265,088	11/1993	Takigawa et al.	370/15
5,267,232	11/1993	Katsube et al.	370/17
5,268,897	12/1993	Komine et al.	370/60
5,271,010	12/1993	Miyake et al.	370/94.1
5,272,697	12/1993	Fraser et al.	370/61
5,274,641	12/1993	Shobatake et al.	370/94.1
5,274,768	12/1993	Traw et al.	395/275
5,280,469	1/1994	Taniguchi et al.	370/13
5,280,470	1/1994	Buhrke et al.	370/13
5,282,201	1/1994	Frank et al.	370/94.1
5,283,788	2/1994	Morita et al.	370/110.1
5,285,446	2/1994	Yonehara	370/60.1
5,287,349	2/1994	Hyodo et al.	370/60.1
5,287,535	2/1994	Sakagawa et al.	370/60
5,289,462	2/1994	Ahmadi et al.	370/60.1
5,289,463	2/1994	Mobasser	370/68.1
5,289,470	2/1994	Chang et al.	370/94.1
5,291,481	3/1994	Doshi et al.	370/79
5,291,482	3/1994	McHarg et al.	370/60
5,295,134	3/1994	Yoshimura et al.	370/16
5,301,055	4/1994	Bagchi et al.	359/139
5,301,155	4/1994	Bagchi et al.	359/139
5,301,184	4/1994	Uriu et al.	370/60
5,301,190	4/1994	Tsukuda et al.	370/66
5,301,193	4/1994	Toyofuku et al.	370/94.1
5,303,232	4/1994	Faulk, Jr.	370/94
5,305,311	4/1994	Lyles	370/60
5,309,431	5/1994	Tominaga et al.	370/60
5,309,438	5/1994	Nakajima	370/94.1
5,311,586	5/1994	Bogart et al.	379/221
5,313,454	5/1994	Bustini et al.	370/13
5,313,458	5/1994	Suzuki	370/56
5,315,586	5/1994	Charvillat	370/60
5,319,638	6/1994	Lin	370/60
5,321,695	6/1994	Proctor et al.	370/60
5,323,389	6/1994	Bitz et al.	370/60.1
5,333,131	7/1994	Tanabe et al.	370/54
5,333,134	7/1994	Ishibashi et al.	370/94.1
5,335,222	8/1994	Kamoi et al.	370/60
5,335,325	8/1994	Frank et al.	395/200
5,339,310	8/1994	Taniguchi	370/60
5,339,317	8/1994	Tanaka et al.	370/85.15
5,339,318	8/1994	Tanaka et al.	370/110.1
5,341,366	8/1994	Soumiya et al.	370/17
5,341,373	8/1994	Ishibashi et al.	370/85.9
5,341,376	8/1994	Yamashita	370/99
5,341,483	8/1994	Frank et al.	395/400
5,345,229	9/1994	Olnowich et al.	340/825.8
5,350,906	9/1994	Brody et al.	235/379
5,355,372	10/1994	Sengupta et al.	370/60
5,357,506	10/1994	Sugawara	370/60
5,357,507	10/1994	Hughes et al.	370/60
5,357,508	10/1994	Le Boudec et al.	370/58.3
5,357,510	10/1994	Norizuki et al.	370/60.1
5,359,600	10/1994	Ueda et al.	370/60.1
5,361,251	11/1994	Aihara et al.	370/60
5,361,372	11/1994	Rege et al.	395/800
5,363,433	11/1994	Isono	379/92
5,363,497	11/1994	Baker et al.	395/425
5,365,514	11/1994	Hershey et al.	370/17
5,367,520	11/1994	Cordell	370/395
5,369,570	11/1994	Parad	364/401
5,371,893	12/1994	Price et al.	395/725
5,373,504	12/1994	Tanaka et al.	370/60.1
5,375,117	12/1994	Morita et al.	370/79
5,377,262	12/1994	Bales et al.	379/220
5,377,327	12/1994	Jain et al.	395/200
5,379,297	1/1995	Glover et al.	370/60.1
5,379,418	1/1995	Shimazaki et al.	395/575
5,390,170	2/1995	Sawant et al.	370/58.1
5,390,174	2/1995	Jugel	370/60
5,390,175	2/1995	Hiller et al.	370/60
5,392,280	2/1995	Zheng	370/353
5,392,402	2/1995	Robrock, II	395/200
5,394,396	2/1995	Yoshimura et al.	370/60
5,394,397	2/1995	Yanagi et al.	370/60.1
5,398,235	3/1995	Tsuzuki et al.	370/16
5,400,337	3/1995	Munter	370/60.1

5,402,415	3/1995	Turner	370/60
5,412,648	5/1995	Fan	370/60
5,414,703	5/1995	Sakaue et al.	370/60
5,418,942	5/1995	Krawchuk et al.	395/600
5,420,858	5/1995	Marshall et al.	370/60.1
5,420,988	5/1995	Elliott	395/275
5,422,879	6/1995	Parsons et al.	370/60
5,425,021	6/1995	Derby et al.	370/54
5,425,026	6/1995	Mori	370/60
5,426,635	6/1995	Mitra et al.	370/60
5,432,713	7/1995	Takeo et al.	364/514
5,432,784	7/1995	Ozveren	370/79
5,432,785	7/1995	Ahmed et al.	370/79
5,432,908	7/1995	Heddes et al.	395/250
5,436,886	7/1995	McGill	370/16
5,436,893	7/1995	Barnett	370/60.1
5,440,547	8/1995	Easki et al.	370/60
5,444,702	8/1995	Burnett et al.	370/60.1
5,446,733	8/1995	Tsuruoka	370/60.1
5,446,737	8/1995	Cidon et al.	370/85.5
5,446,738	8/1995	Kim et al.	370/94.2
5,448,559	9/1995	Hayter et al.	370/60.1
5,448,621	9/1995	Knudsen	379/58
5,450,406	9/1995	Esaki et al.	370/60.1
5,452,296	9/1995	Shimizu	370/60.1
5,454,299	10/1995	Thessin et al.	370/62
5,455,820	10/1995	Yamada	370/17
5,455,825	10/1995	Lauer et al.	370/60
5,457,687	10/1995	Newman	370/85.3
5,459,743	10/1995	Fukuda et al.	371/67.1
5,461,611	10/1995	Drake, Jr. et al.	370/54
5,463,620	10/1995	Sriram	370/60
5,463,629	10/1995	Ko	370/110.1
5,463,775	10/1995	DeWitt et al.	395/184.01
5,465,331	11/1995	Yang et al.	395/200
5,465,365	11/1995	Winterbottom	395/600
5,469,003	11/1995	Kean	326/39
5,473,608	12/1995	Gagne et al.	370/85.13
5,475,679	12/1995	Munter	370/58.2
5,479,401	12/1995	Bitz et al.	370/60.1
5,479,402	12/1995	Hata et al.	370/60.1
5,483,526	1/1996	Ben-Nun et al.	370/60.1
5,485,453	1/1996	Wahlman et al.	370/16
5,485,455	1/1996	Dobbins et al.	370/60
5,487,063	1/1996	Kakuma et al.	370/56
5,488,606	1/1996	Kakuma et al.	370/16
5,491,691	2/1996	Shtayer et al.	370/61
5,491,694	2/1996	Oliver et al.	370/85.4
5,493,566	2/1996	Ljungberg et al.	370/60
5,497,369	3/1996	Wainwright	370/60
5,499,238	3/1996	Shon	370/60.2
5,504,741	4/1996	Yamanaka et al.	370/58.2
5,504,742	4/1996	Kakuma et al.	370/60.1
5,506,834	4/1996	Sekihata et al.	370/17
5,506,839	4/1996	Hatta	370/60
5,506,956	4/1996	Cohen	395/182.04
5,509,001	4/1996	Tachibana et al.	370/17
5,509,007	4/1996	Takashima et al.	370/60.1
5,513,134	4/1996	Cooperman et al.	365/49
5,513,178	4/1996	Tanaka	370/58.2
5,513,180	4/1996	Miyake et al.	370/60.1
5,515,359	5/1996	Zheng	370/13
5,517,495	5/1996	Lund et al.	370/399
5,519,690	5/1996	Suzuka et al.	370/17
5,521,905	5/1996	Oda et al.	370/17
5,521,915	5/1996	Dieudonne et al.	370/60.1
5,521,916	5/1996	Choudhury et al.	370/60.1
5,521,917	5/1996	Watanabe et al.	370/1
5,521,923	5/1996	Willmann et al.	370/94.1
5,523,999	6/1996	Takano et al.	370/58.2
5,524,113	6/1996	Gaddis	370/60.1
5,526,344	6/1996	Diaz et al.	370/16
5,528,588	6/1996	Bennett et al.	370/60
5,528,590	6/1996	Iidaka et al.	370/60.1
5,528,591	6/1996	Lauer	370/79
5,530,695	6/1996	Digne et al.	370/17
5,533,009	7/1996	Chen	370/17
5,533,020	7/1996	Byrn et al.	370/60.1
5,535,196	7/1996	Aihara et al.	370/60
5,535,197	7/1996	Cotton	370/60
5,537,394	7/1996	Abe et al.	370/17
5,541,912	7/1996	Choudbury et al.	370/17
5,544,168	8/1996	Jeffrey et al.	370/60.1
5,544,169	8/1996	Norizuki et al.	370/60.1
5,544,170	8/1996	Kasahara	370/84
5,546,389	8/1996	Wippenbeck et al.	370/60
5,546,391	8/1996	Hochschild et al.	370/60
5,546,392	8/1996	Boal et al.	370/60.1
5,550,821	8/1996	Akiyoshi	370/60.1
5,550,823	8/1996	Irie et al.	370/60.1
5,553,057	9/1996	Nakayama	370/13
5,553,068	9/1996	Aso et al.	370/60
5,555,243	9/1996	Kakuma et al.	370/58.2
5,555,265	9/1996	Kakuma et al.	370/60
5,557,607	9/1996	Holden	370/58.2
5,568,479	10/1996	Watanabe et al.	370/60.1
5,570,361	10/1996	Norizuki et al.	370/60.1
5,570,362	10/1996	Nishimura	370/60.1
5,572,522	11/1996	Calamvokis et al.	370/60.1
5,577,032	11/1996	Sone et al.	370/58.3
5,577,035	11/1996	Hayter et al.	370/60
5,583,857	12/1996	Soumiya et al.	370/233
5,583,858	12/1996	Hanaoka	370/392
5,583,861	12/1996	Holden	370/395
5,590,132	12/1996	Ishibashi et al.	370/236
5,602,829	2/1997	Nie et al.	370/235
5,610,913	3/1997	Tomonaga et al.	370/219
5,610,914	3/1997	Yamada	370/395
5,623,405	4/1997	Isono	395/230
5,625,846	4/1997	Kobayakawa et al.	395/872
5,633,861	5/1997	Hanson et al.	370/232
5,640,389	6/1997	Masaki et al.	370/418

OTHER PUBLICATIONS

Sriram, K., "Methodologies for Bandwidth Allocation, Transmission Scheduling, and Congestion Avoidance in Broadband ATM Networks," Globecom '92: IEEE Global Telecommunications Conference, pp. 1545–1551, Feb. 1992.

Zukowski, C., et al., "VLSI Implementations of ATM Buffer Management," Communications, IEEE International Conference, pp. 716–720, 1991.

Endo, N., et al., "Traffic Characteristics Evaluation of a Shared Buffer ATM Switch," IEEE Global Telecommunications Conference, pp. 1913–1918, 1990.

Guerin, R., et al., "Overflow Analysis for Finite Waiting Room Systems," IEEE Transactions on Communications, vol. 38 No. 9, pp. 1569–1577, Sep. 1990.

An Ascom Timeplex White Paper, *Meeting Critical Requirements with Scalable Enterprise Networking Solutions Based on a Unified ATM Foundation,* pp. 1–12, Apr. 1994.–Apr. 1995?.

Douglas H. Hunt, *ATM Traffic Management—Another Perspective,* Business Communications Review, Jul. 1994.

Richard Bubenik et al., *Leaf Initiated Join Extensions,* Technical Committee, Signalling Subworking Group, ATM Forum/94–0325R1, Jul. 1, 1994.

Douglas H. Hunt et al., *Flow Controlled Virtual Connections Proposal for ATM Traffic Management (Revision R2)*, Traffic Management Subworking Group, ATM_Forum/94–0632R2, Aug. 1994.

Flavio Bonomi et al., *The Rate–Based Flow Control Framework for the Available Bit Rate ATM Service,* IEEE Network, Mar./Apr. 1995, pp. 25–39.

R. Jain, *Myths About Congestion Management in High Speed Networks,* Internetworking Research and Experience, vol. 3, 101–113 (1992).

Douglas H. Hunt et al., *Credit–Based FCVC Proposal for ATM Traffic Management (Revision R1)* ATM Forum Technical Committee Traffic Management Subworking Group, ATM_Forum/94–0168R1, Apr. 28, 1994.

Douglas H. Hunt et al., *Action Item Status for Credit–Based FCVC Proposal,* ATM Forum Technical Committee Traffic Management Subworking Group, ATM_Forum/94–0439, Apr. 28, 1994.

Timothy P. Donahue et al., *Arguments in Favor of Continuing Phase 1 as the Initial ATM Forum P–NNI Routing Protocol Implementation,* ATM Forum Technical Committee, ATM Forum/94–0460, Apr. 28, 1994.

Richard Bubenick et al., *Leaf Initiated Join Extensions,* Technical Committee, Signalling Subworking Group, ATM Forum/94–0325, Apr. 28, 1994.

Rob Coltun et al., *PRP: A P–NNI Routing Protocol Proposal,* ATM Forum Technical Committee, ATM_Forum/94–0492, Apr. 28, 1994.

Richard Bubenik et al., *Leaf Initiated Join Extensions,* ATM Forum Technical Committee, Signalling Subworking Group, ATM Forum 94–0325, Apr. 28, 1994.

Richard Bubenik et al., *Requirements For Phase 2 Signaling Protocol,* ATM Forum Technical Committee, Signalling Subworking Group, ATM Forum 94–1078, Jan. 1, 1994.

H.T. Kung and K. Chang, *Receiver–Oriented Adaptive Buffer Allocation in Credit–Based Flow Control for ATM Networks,* Proceedings of INFOCOM '95, Apr. 2–6, 1995, pp. 1–14.

H.T. Kung et al., *Credit–Based Flow Control for ATM Networks: Credit Update Protocol, Adaptive Credit Allocation, and Statistical Multiplexing,* Proceedings of ACM SIGCOMM '94 Symposium on Communications Architectures, Protocols and Applications, Aug. 31–Sep. 2, 1994, pp. 1–14.

SITA, ATM RFP: C–Overall Technical Requirements, Sep. 1994.

Head of Line Arbitration in ATM Switches With Input–Output Buffering and Backpressure Control. By Hosein F. Badran and H. T. Mouftah, *Globecom* '91, pp. 0347–0351.

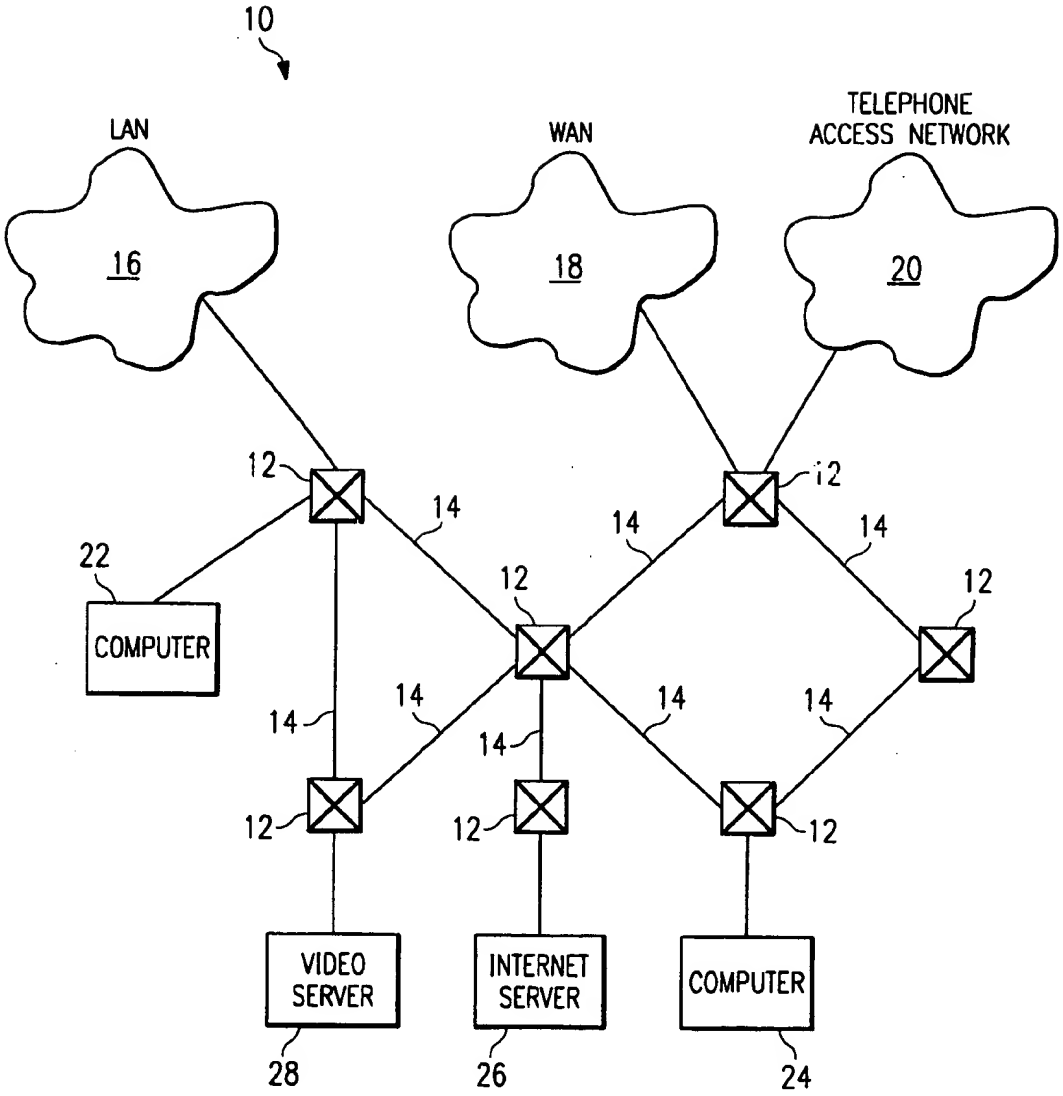
10
LAN
WAN
TELEPHONE ACCESS NETWORK
16
18
20
12
14
22
COMPUTER
VIDEO SERVER
INTERNET SERVER
COMPUTER
28
26
24

FIG. 1

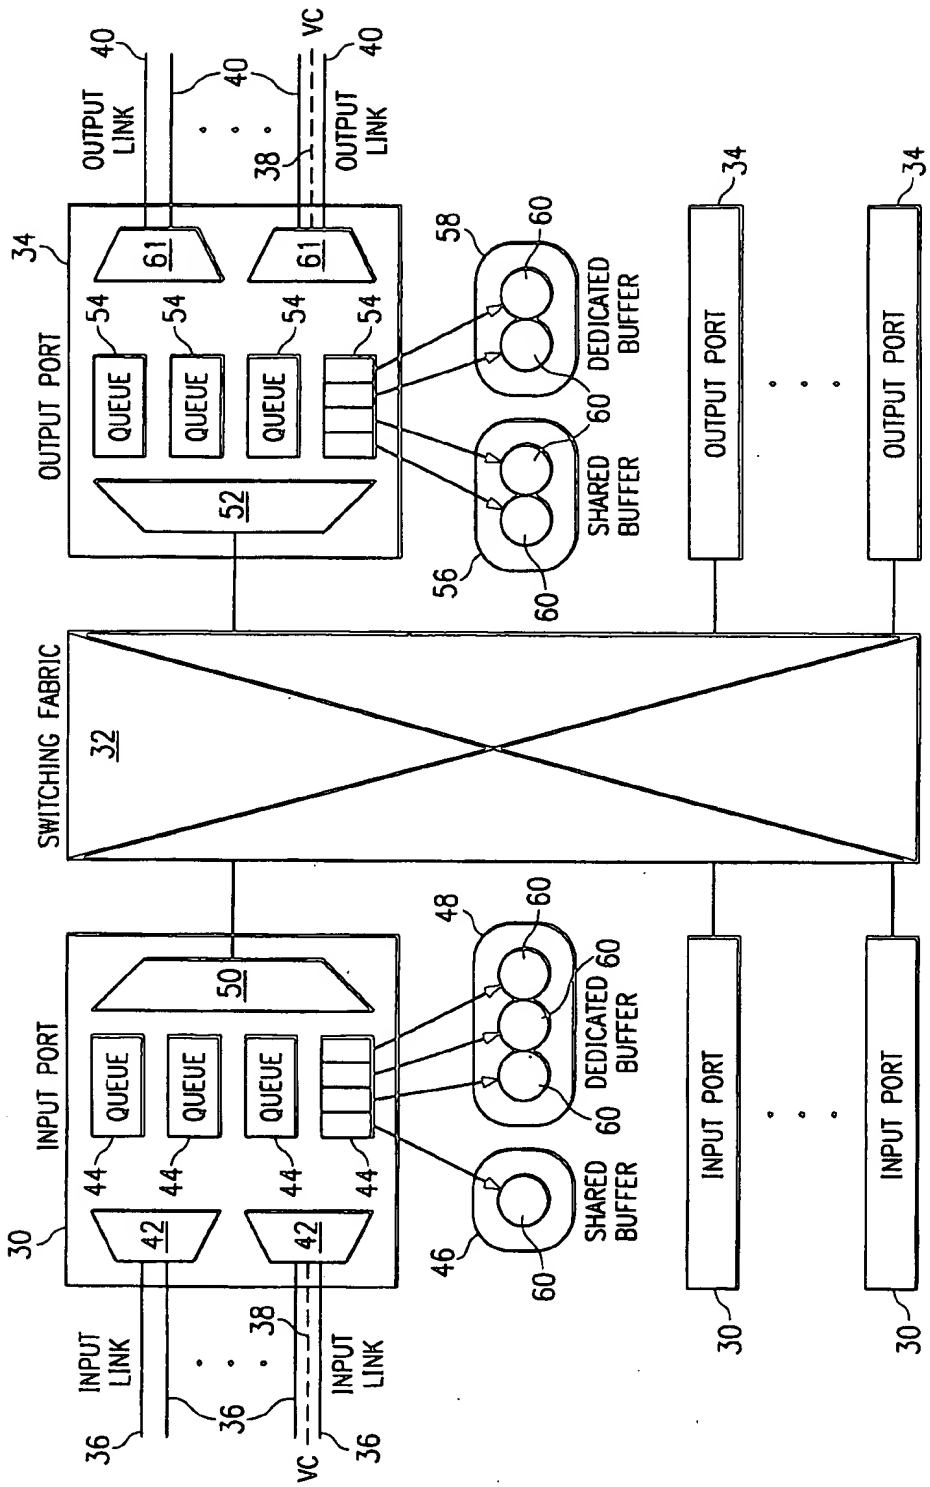
OUTPUT LINK
40
VC
38
OUTPUT PORT
34
61
QUEUE
54
52
58
60
DEDICATED BUFFER
SHARED BUFFER
56
SWITCHING FABRIC
32
INPUT PORT
30
42
44
50
48
46
INPUT LINK
36
12

FIG. 2

METHOD AND SYSTEM FOR CONTROLLING NETWORK SERVICE PARAMETERS IN A CELL BASED COMMUNICATIONS NETWORK

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional application Ser. No. 60/001,498, filed Jul. 19, 1995.

TECHNICAL FIELD OF THE INVENTION

This invention relates generally to communications networks and more particularly to a method and system for controlling network service parameters in cell based communications networks.

BACKGROUND OF THE INVENTION

Cell based communications networks such as asynchronous transfer ("ATM") networks are used to transfer audio, video, and data. ATM networks transfer data by routing data units such as ATM cells from a source to a destination through switches. Switches include input/output ports through which ATM cells are received and transmitted.

Existing ATM switches can be divided into several classes. One type of switch uses one large queue as a buffer and controls the access of individual virtual connections to bandwidth on a statistical basis. As a general matter, each virtual connection through this type of switch shares one large buffer and each virtual connection statistically shares the available bandwidth of the switch and/or on an output link connected to the switch. This approach suffers from several disadvantages. Depending upon the type of traffic traveling through an ATM network, one may want to control certain network service parameters associated with that traffic type. Examples of parameters that may be desirably controlled by an ATM network include a connection's delay variation, its end-to-end delay, its maximum data loss, its access to guaranteed bandwidth, its access and priority of access to shared bandwidth, and its delivery priority. Unfortunately, ATM switches employing the above approach cannot create quality of service guarantees on a per connection basis or even provide quality of service (QOS) guarantees on a per traffic type basis. Instead, quality of service guarantees can only be statistical in nature. In addition, head of line blocking in this approach can block connections behind the blocked connection.

A second approach employs a plurality of queues in each switch and allocates switch bandwidth statistically. Further refinements of this approach may assign particular connections to a particular queue based upon a characteristic of the connection such as the traffic type of the connection or the group from which the connection originated. In general, this approach is characterized by sharing each of a plurality of queues by a plurality of connections and sharing the switch bandwidth and/or output link bandwidth among all connections. This approach also has several disadvantages. Quality of service based upon the above parameters can only be guaranteed on a per queue basis and the guarantees that can be made for low priority traffic are not strong guarantees. In addition, the quality of service guarantees for multiple queues can only be made on a statistical basis rather than on a per connection basis. This approach also suffers from head of line blocking in each individual queue. Thus, certain connections may block other connections behind it.

A third approach uses a plurality of queues and assigns a plurality of connections to each queue and allocates either

switch bandwidth and/or link bandwidth on a per queue basis. This approach also suffers from several disadvantages. Again, the quality of network service can only be guaranteed on a per queue basis rather than on a per connection basis. Also, unused allocated switch bandwidth and/or link bandwidth cannot be used by another queue. Accordingly, this bandwidth remains unused, resulting in an inefficient use of system resources.

SUMMARY OF THE INVENTION

The invention comprises a method and system for controlling network service parameters in a cell based communications network. In accordance with a method of the invention, a plurality of input signals are received on input ports of a communications device wherein at least one input signal comprises at least one virtual connection. The communications device is part of a cell based communications network wherein each virtual connection comprises a series of data cells comprising a header portion and a data portion. A first virtual connection is assigned to a first queue. The first queue is associated with a first buffer space wherein the first buffer space is dedicated to the first queue. The first queue is also associated with a second buffer space wherein the second buffer space is shared between the first queue and other queues. Data associated with the first virtual connection is queued in the first queue.

The invention allows quality of service guarantees based upon network service parameters to be made on a per connection basis. Thus, the invention can be used for efficiently transporting disparate types of traffic across a cell based communications network. For constant bit rate (CBR) services, the invention allows delay variation, end-to-end delay and bandwidth to be guaranteed on a per virtual connection basis. In addition, the invention allows the guarantee of "no cell loss" for a constant bit rate traffic other than link bit errors. The term "no cell loss" refers to preventing cell loss caused by network congestion due to buffer overflow.

For variable bit rate services, a portion of switch and link resources may be allocated to a variable bit rate (VBR) virtual connection and other switch and link resources may be dynamically shared among that particular virtual connection and other traffic on the network. For the portion of services that are allocated, the delay variation, end-to-end delay, bandwidth, and maximum data loss can be guaranteed for a variable bit rate virtual connection. For bandwidth above the allocated portion, maximum data loss can be guaranteed on a statistical basis. In addition, priority access to dynamic bandwidth and delivery priority can be guaranteed for a particular variable bit rate virtual connection.

For an available bit rate (ABR) virtual connection, the invention can guarantee no cell loss and priority access to dynamic bandwidth either in a switch, on a link, or both. Like variable bit rate services, a portion of switch and/or link resources may be allocated to an available bit rate virtual connection while other resources may be shared with other virtual connections. For the allocated portion of a virtual connection, the bandwidth of that connection can be guaranteed.

For unspecified bit rate services, the invention allows a guarantee of priority access to dynamic bandwidth and of delivery priority.

Accordingly, the invention allows quality of service guarantees to be made for disparate types of traffic that may be transported using a communications network employing the present invention. The invention avoids the shortcomings of

the prior art by applying queuing and bandwidth management mechanisms to individual connections as well as to a group of connections. Without the per virtual connection queuing mechanism of the invention, resources may not be precisely allocated to individual flows so that quality of service guarantees are enduring for the life of a connection. The invention allows quality of service guarantees to be maintained regardless of how many other virtual connections are added or deleted during the duration of the connection.

A communications network constructed in accordance with the invention may also achieve 100% network occupancy rates on a link-by-link basis while maintaining constant bit rate bandwidth and delay guarantees and available bit rate quality of service guarantees in an integrated services environment. The invention thus allows efficient use of network resources by devices connected to the communications network, that have varying traffic patterns and quality of service needs.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawings in which:

FIG. **1** illustrates a communications network constructed in accordance with the teachings of the invention; and

FIG. **2** illustrates a first embodiment of a communications device constructed in accordance with the teachings of the invention.

DETAILED DESCRIPTION OF THE INVENTION

The preferred embodiment of the present invention and its advantages are best understood by referring to FIGS. **1** and **2** of the drawings, like numerals being used for like and corresponding parts of the various drawings.

FIG. **1** illustrates a communications network **10** constructed in accordance with the teachings of the invention. Communications network **10** is a cell based network and in this embodiment comprises an asynchronous transfer mode (ATM) network. Communications network **10** could be another type of network without departing from the scope of the invention. Communications network **10** comprises a series of communications devices **12** (also referred to as nodes) connected by a series of communications links **14**. In this embodiment, each communications device **12** comprises an ATM switch. Other types of switches could be used without departing from the scope of the invention. Each communications link **14** may comprise any type of communications link, such as a SONET link.

A number of communications devices and networks may be connected to communications network **10**. In this embodiment, communications network **10** is connected to local area network **16**, wide area network **18**, telephone access network **20**, first computer **22**, second computer **24**, Internet server **26** and video server **28**. Other types of communications devices (or networks) could also be connected to communications network **10** and some or all of the illustrated devices (or networks) could be omitted without departing from the scope of the invention. Each of these communications devices (or networks) communicates with communications network **10** using ATM cells. These devices (or networks) may use a communications protocol other than ATM for internal communications. For such devices, a mechanism is provided to translate the communications protocol to ATM and from ATM to the communications protocol for the particular device or network. This mechanism may be provided either inside of one of the communications devices **12** or inside the particular device or network connected to communications network **10**.

When a first device connected to communications network **10** desires to communicate with a second device connected to communications network **10**, it sends data over communications network **10**. The data is sent from the first device over one or more communication links **14** and through one or more communications devices **12** to the second device, thus forming a virtual connection between the first device and the second device. A communications device **12** switches numerous virtual connections through communications device **12** and each communications link **14** carries numerous virtual connections between two communications devices **12**.

Communications network **10** is an integrated services network. Thus, communications network **10** can be used to transport varying types of ATM services. Examples of ATM services that may be transported through communications network **10** include constant bit rate services, variable bit rate services, available bit rate services, and unspecified bit rate services. The invention allows efficient transport of each of these traffic types and allows a certain quality of service guarantees to be made regarding each traffic type or virtual connection.

In accordance with one aspect of the invention, a unique queue is assigned to each virtual connection passing through a communications device **12**. A specific embodiment of per virtual connection queuing will be described in connection with FIG. **2** below. In the embodiment illustrated in FIG. **2**, both an input queue and an output queue are associated with each virtual connection. These queues are uniquely assigned to a particular virtual connection. Thus, queues are not shared among virtual connections. Some switches may employ only an input queue, only an output queue, or queues in addition to an input and output queue. The teachings of the invention can be applied to any queue associated with a communications device **12** by assigning queues on a per connection basis. As will be discussed more fully below, per virtual connection queuing allows queuing and bandwidth management mechanisms to be applied to individual connections rather than a group of connections. Queuing management mechanisms include control over buffering of virtual connections. Per virtual connection queuing allows precise allocation of network resources to individual virtual connections so that quality of service guarantees are enduring for the life of the virtual connection. These guarantees may be maintained regardless of how many other connections are added or deleted.

In operation, a plurality of input signals are received on input ports of communications device **12** wherein at least one input signal includes at least one virtual connection. In this embodiment, a unique input queue and a unique output queue are associated with each virtual connection. Queues in this embodiment are logical structures implemented using a linked list of pointers. Each queue may be associated with one or more buffer pools. In this embodiment, each queue may be associated with two buffer pools—a shared buffer pool and a dedicated buffer pool. Data associated with the first virtual connection is queued in the unique input queue for transport through a switching fabric connected with the input port. Data that is passed through the switching fabric is queued in the unique output queue associated with the virtual connection for transport over an output communica-

tions link. As will be described more fully below, another aspect of the invention is the method of implementing per virtual connection queuing.

A second aspect of the invention involves the use of communications network **10** in emulating a circuit type connection. Devices using a constant bit rate mode of communication frequently need to be treated as if there was a physical connection between the devices. Accordingly, ATM networks sometimes include resources to aid in emulating circuit connections so as to provide adequate service between devices communicating using a constant bit rate service. Communications network **10** of the present invention may be used in the emulation of circuit connections while guaranteeing certain network quality of service parameters. For each constant bit rate virtual connection, the present invention allows quality of service to be guaranteed for delay variation and end-to-end delay on a per virtual connection basis. In addition, bandwidth can be guaranteed for a constant bit rate virtual connection. Also, constant bit rate services may be transported through communications network **10** with no cell loss.

Although a more specific mechanism for implementing circuit connection emulation will be described in connection with FIG. **2** below, the above-described per virtual connection queuing mechanism may be used to emulate a circuit connection. For example, first computer **22** of FIG. **1** may desire to communicate using a constant bit rate service with second computer **24**. To achieve circuit connection emulation, queues and bandwidth through each communications device **12** through which the virtual connection between computer **22** and computer **24** passes are allocated to the virtual connection for the duration of the communication signal between first computer **22** and second computer **24**. The queue that is allocated has access to an allocated amount of buffer space for the duration of the connection. In this embodiment, allocated buffer space is implemented by assigning a counter limit to the queue. This counter limit represents the number of cell locations in a buffer pool that may be used by the particular queue at any one time. All queues that have allocated buffer space may use any available physical location in the shared pool up to the counter limit. In addition, an amount of output link bandwidth is allocated to the communications signal over each communications link **14** of communications network **10** over which the virtual connection travels between first computer **22** and second computer **24**. By allocating network resources and dedicating those resources to a particular virtual connection, a circuit connection may be emulated through communications network **10** with quality of service guarantees on a per connection basis.

In this specific embodiment, communications device **12** has both an input queue and an output queue assigned to a constant bit rate virtual connection. These queues are created when a new virtual connection is established. In the current embodiment, 16,000 connections per input port or output port may be processed. In addition, bandwidth through the switching fabric of communications device **12** and bandwidth on the output link of communications device **12** are allocated to the constant bit rate virtual connection. Where communications device **12** is configured differently, circuit connections may be emulated by allocating any queuing and bandwidth resources on a per virtual connection basis in accordance with the invention.

FIG. **2** illustrates a first embodiment of a communications device **12** constructed in accordance with the teachings of the invention. FIG. **2** comprises a plurality of input ports **30** coupled to a switching fabric **32** which is, in turn, coupled to a plurality of output ports **34**. The term "coupled" refers to a logical connection between the input ports **30**, switching fabric **32** and output ports **34**. Various additional circuitry may appear between these elements, but they are still logically coupled to one another. A larger or smaller number of input ports **30** and/or output ports **34** could be used without departing from the scope of the invention.

Each input port **30** is connected to one or more input links **36**. Each input link **36** in turn comprises one or more virtual connections **38**. Each output port **34** is connected to one or more output links **40**. Each output link **40** in turn comprises one or more virtual connections **38**.

Each input port **30** comprises one or more demultiplexers **42**, one or more queues **44**, and a multiplexer **50**. Demultiplexers **42** and multiplexer **50** are not physical devices. Instead, the switching of traffic through input port **30** functionally implements demultiplexing and multiplexing functions.

When a virtual connection **38** is received over an input link **36** the cells associated with virtual connection **38** are demultiplexed through a demultiplexer **42** to one of the queues **44**. In this embodiment, a unique queue **44** is assigned to each virtual connection **38**. In other words, queue **44** contains the cells of one and only one input link **38**. Although this embodiment employs per virtual connection queuing, the buffer and bandwidth mechanisms of the invention may be used in communications devices **12** that assign multiple virtual connections to a single queue **44**.

In this embodiment, queue **44** comprises a linked list of pointers. These pointers point to the location in a buffer where a cell associated with a particular queue pointer is stored. A third aspect of the invention is the queuing mechanism used for queues **44**. Each queue **44** may be associated with multiple buffer pools. In accordance with the queuing mechanism of this embodiment, queue **44** is associated with a particular virtual connection **38** and may be assigned dedicated buffer pool space, shared buffer pool space, or a combination of dedicated and shared buffer pool space depending upon the traffic type of the connection.

Referring to FIG. **2**, one of the input queues **44** has been assigned space in shared buffer **46** and dedicated buffer **48**. Space in dedicated buffer **48** is uniquely dedicated to a particular queue **44**. Thus, space in dedicated buffer **48** associated with a queue **44** may be used by that queue **44** and only that queue **44**. Space in shared buffer **46** is shared among several queues **44**. Ordinarily, a constant bit rate virtual connection **38** will only be assigned space in a dedicated buffer **48**, while a variable bit rate virtual connection **38** and available bit rate virtual connection **38**, or an unspecified bit rate virtual connection **38** may be assigned dedicated space in a dedicated buffer **48** and shared space in a shared buffer **46**. In this embodiment, because a particular virtual connection **38** is uniquely assigned a queue **44**, space in a dedicated buffer **48** and/or in a shared buffer **46** may be assigned specifically to a virtual connection. In an embodiment where queue **44** is shared among several virtual connections **38**, space in dedicated buffer **48** and/or shared buffer **46** may be assigned to the queue **44** as a whole.

Dedicated buffer **48** comprises a plurality of buffer locations **60** wherein each buffer location **60** holds one ATM cell. The size of dedicated buffer **48** may be dynamically adjusted depending upon the number of virtual connections accessing dedicated buffer **48**. When a queue **44** is assigned space in dedicated buffer **48**, the queue **44** is assigned a fixed number of buffer locations **60** that the queue **44** may have access to at any one time. The specific physical position of those

buffer locations **60** are not assigned to any specific queue **44**. Rather, dedicated buffer **48** acts as a pool of buffer space. Any queue **44** that is assigned space in dedicated buffer **48** may use any otherwise unused physical location within it, as long as that queue **44** does not use more than the fixed number of buffer locations **60** that it is assigned. The queue **44** is guaranteed to have access to the fixed number of buffer locations **60** assigned to it at all times.

Similarly, shared buffer **46** acts as a pool of buffer space that can be shared in a similar manner to dedicated buffer **48**. Each queue **44** is associated with a limit specifying the maximum number of locations in shared buffer **46** that the queue **44** may use at any one time. Here, however, each queue **44** is not guaranteed to have access to any fixed number of buffer locations **60**. Instead, all queues **44** assigned space in shared buffer **46** share all buffer locations **60**. When all of the buffer locations **60** in shared buffer **46** are in use, a queue **44** may not have access to another buffer location **60** even if it is not using the maximum number of buffer locations **60** that could be assigned to it. The size of shared buffer **46** may also be adjusted dynamically. A connection could be assigned to multiple shared buffers **46** or connections could be assigned to different shared buffers **46** based upon their traffic types.

Each of the queues **44** are connected to multiplexer **50**. Multiplexer **50** controls the flow of cells between each input port **30** and switching fabric **32**. As discussed above, multiplexer **50** is not a physical element, but is functionally implemented by the operation of input port **30**.

In this embodiment, multiplexer **50** is used to assign bandwidth through switching fabric **32** to each of the queues **44** connected to multiplexer **50**. This embodiment of the invention employs a unique bandwidth assignment method. A portion of the bandwidth through switching fabric **32** may be uniquely allocated to a particular queue **44** such that a fixed amount of bandwidth through switching fabric **32** is guaranteed to that queue **44**. In addition, an amount of bandwidth through switching fabric **32** may be dynamically shared between two or more queues **44**. Thus, an amount of uniquely allocated bandwidth through switching fabric **32** may be assigned to a particular queue **44**. In addition, that particular queue **44** may or may not share in the available shared dynamic bandwidth through the switching fabric **32**. In this embodiment, because queue **44** is associated with only one virtual connection, a bandwidth assignment is equivalent to assigning an amount of uniquely allocated bandwidth to a particular virtual connection. The same virtual connection may or may not share dynamic bandwidth. This aspect of the invention could also be used in a communications device **12** that assigned multiple virtual connections to a single queue **44**. In such an embodiment, uniquely allocated bandwidth and/or shared dynamic bandwidth could be used by a queue **44**.

Space in shared buffer **46** and/or dynamically shared bandwidth through switching fabric **32** may be shared by all queues **44** or may be divided into a series of pools shared by specific groups of the queues **44**. For example, dynamic buffer space and/or bandwidth could be divided into pools associated with particular traffic types. Variable bit rate virtual connections, available bit rate virtual connections, and/or unspecified bit rate virtual connections could share buffer space and/or dynamic bandwidth in a pooled manner. A mechanism for achieving this type of sharing will be described more fully below.

If switching fabric **32** comprises a switching matrix, bandwidth and/or buffers in switching fabric **32** could be shared using the mechanisms described above wherein bandwidth and buffers may be either allocated or shared.

Each output port **34** comprises a demultiplexer **52**, one or more queues **54**, and one or more multiplexers **61**. Each output port **34** is connected to one or more output links **40** and each output link **40** comprises one or more virtual connections **38**. As was the case with input port **30**, demultiplexer **52** and multiplexer **61** are not physical elements of output port **34**, but are functionally implemented therein.

Demultiplexer **52** sends cells received from switching fabric **32** to the appropriate output queue **54**. In this embodiment, each output queue **54** is associated with a unique virtual connection. In other words, each output queue **54** is used to queue the cells associated with one and only one virtual connection **38**.

The queues **54** are similar to the queues **44**. Each queue **54** may be associated with multiple buffer pools. As was the case with the queues **44** of input port **30**, queues **54** of output port **34** comprise a linked list of pointers. Each pointer points to a cell location **60** in either a shared buffer **56** or dedicated buffer **58**. Each queue **54** may be assigned an amount of uniquely dedicated space in a dedicated buffer **58** and/or an amount of shared space in a shared buffer **56**. Accordingly, output queues **54** employ a buffering mechanism similar to that of input queues **44**.

Each multiplexer **61** is connected to each of the output queues **54**. As was the case with the multiplexer **50** and input port **30**, the multiplexer **61** of output **34** are not physical elements of output **34**, but output port **34** implements a multiplexing function. Each multiplexer **61** controls the access of the output queues **54** to bandwidth on its associated output link **40**. Bandwidth on each output link **40** may be allocated as described above in connection with input port **30** for the input queues **44**. Thus, each output queue **54** may be assigned an amount of uniquely allocated bandwidth which that particular output queue **54** is guaranteed to have access to if needed. That output queue **54** may also be assigned an amount of dynamically shared bandwidth that is shared among multiple output queues **54**. Each output queue **54** may be assigned an amount of dynamic bandwidth, an amount of shared bandwidth, or a combination of both on a particular output link **40**. In this embodiment, because each output queue **54** is only associated with a single virtual connection **38**, this means that bandwidth may be allocated uniquely and/or dynamically shared on a per virtual connection basis. In an embodiment where an output queue **54** is shared among multiple virtual connections **38**, bandwidth may be uniquely allocated or dynamically shared on a per queue basis using this mechanism.

As described above in connection with input port **30**, space in a shared buffer **56** and/or dynamically shared bandwidth on an output link **40** may be shared among all virtual connections **30** or divided up and shared in pools corresponding to particular group of virtual connections such as traffic types. Again, this mechanism will be described more fully below.

Sometimes, a particular virtual connection **38** will be allocated more bandwidth through the switching fabric **32** than is needed by that virtual connection. In other words, the amount of allocated bandwidth through the switching fabric **32** may be uniquely allocated to a virtual connection having a bandwidth smaller than the amount of allocated bandwidth. Allocating additional bandwidth reduces delay through communications device **12**. In the present invention, bandwidth that is allocated to a particular virtual connection **38** through the switching fabric **32** that is unused by that

particular virtual connection **38** becomes part of the switch's dynamic bandwidth. This dynamic bandwidth is shared between virtual connections using dynamic bandwidth services. This feature of the invention allows more efficient use of the resources of communications network **10**.

Access to shared buffer space may also be prioritized. Each virtual connection may have a priority assigned to it. Each buffer pool has priority thresholds. As the available buffers in the buffer pool decrease below the thresholds, virtual connections with corresponding priorities are denied further access to the buffer pool.

As described above, connections can have access to shared buffers and bandwidth. Levels of control can be added to manage both the order and the amount of the shared resources that connections can use. The highest level of control is to ensure that traffic types receive a guaranteed minimum assignment of the shared resource. This is accomplished at the output link by measuring bandwidth usage per traffic type, and granting each traffic type bandwidth until it has received the desired minimum amount. For buffers, this is accomplished by having a buffer pool per traffic type. The order in which the traffic types receive bandwidth could be a programmable policy. In this embodiment, the order is fixed. Quality of service levels can be determined for connections sharing the minimum bandwidth and buffer guarantee for a traffic type, assuming that the connections' input characteristics can be accurately described. In this embodiment, the connections of the same traffic type are multiplexed together using a linked list data structure at both bandwidth contention points—i.e. the switch port and link. This multiplexing point is given the minimum bandwidth guarantee. An identifier in the connection's queue descriptor (s) is used to associate the connections with a multiplexing point.

Access to the traffic type's shared resources is also ordered. That is, connections within a traffic type get access to bandwidth and buffers using a priority mechanism. Again, this prioritization occurs at the resource contention points, so that there are traffic type priorities at both the input port processor and the output port processor. Input buffer pools have one or more thresholds, above which certain connections within a traffic type are denied buffers. There are one or more multiplexers, i.e. lists, per traffic type, where the multiplexers are assigned, in priority order, bandwidth granted to the traffic type at the output port. This means that low priority connections get bandwidth only after the high priority connections do not have any data to transmit. An identifier in the connection's queue descriptor(s) is used to associate the connections with a priority.

Note that more levels of access control to shared resources can be added in the switch. For example, the lowest level of control occurs on a multipoint-to-point connection. Bandwidth is distributed to the multiple sources of a single multipoint-to-point connection by the use of two additional multiplexers: (1) a list of queues at the input port that distributes the bandwidth between multiple sources merging at the input port, and (2) an arbitration mechanism in the switch fabric that distributes bandwidth between sources on multiple input ports.

Although the present invention has been described in detail, it should be understood that various changes, substitutions, and alterations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.

What is claimed is:

1. A method for controlling network service parameters in a cell based communications network, comprising:

receiving a plurality of input signals on input ports of a communications device, at least one input signal comprising at least one virtual connection, wherein the communications device is part of the cell based communications network and wherein each virtual connection comprises a series of data cells comprising a header portion and a data portion;

assigning a first virtual connection to a first queue;

associating the first queue with a first buffer space, wherein the first buffer space is dedicated to the first queue;

associating the first queue with a second buffer space, wherein the second buffer space is shared between the first queue and other queues;

queueing data associated with the first virtual connection in the first queue.

2. The method of claim **1**, wherein the data cells comprise asynchronous transfer mode cells.

3. The method of claim **1** wherein the first queue comprises an input queue.

4. The method of claim **3**, further comprising:

assigning the first virtual connection to a second queue;

associating the second queue with a third buffer space, wherein the third buffer space is dedicated to the second queue;

associating the second queue with a fourth buffer space, wherein the fourth buffer space is shared between the second queue and other queues;

queueing data associated with the first virtual connection in the second queue wherein the second queue comprises an output queue.

5. The method of claim **4**, further comprising:

controlling the flow of the first virtual connection from an input port of the communications device to a switching fabric associated with the communications device by utilizing the first queue associated with the first virtual connection.

6. The method of claim **5**, further comprising:

controlling the flow of the first virtual connection from a switching fabric associated with the communications device to an output port of the communications device by utilizing the second queue associated with the first virtual connection; and

wherein the first queue is uniquely associated with the first virtual connection and wherein the second queue is uniquely associated with the first virtual connection.

7. The method of claim **6**, wherein the data cells comprise asynchronous transfer mode cells.

8. The method of claim **6**, further comprising:

assigning a first quantity of bandwidth through the switching fabric to the first queue, wherein the first quantity of bandwidth is allocated to the first queue;

assigning a second quantity of bandwidth through the switching fabric to the first queue, wherein the second quantity of bandwidth is dynamically shared between the first queue and other queues; and

passing cells associated with the first virtual connection through the switching fabric using the first quantity of bandwidth and at least a part of the second quantity of bandwidth.

9. The method of claim **6**, further comprising:

assigning a third quantity of bandwidth on an output link of the communications device to the second queue, wherein the third quantity of bandwidth is allocated to the second queue;

assigning a fourth quantity of bandwidth on the output link to the second queue, wherein the fourth quantity of bandwidth is dynamically shared between the second queue and other queues; and

passing cells associated with the first virtual connection to the output link using the third quantity of bandwidth and at least a part of the fourth quantity of bandwidth.

10. The method of claim 9, further comprising:

assigning a first quantity of bandwidth through the switching fabric to the first queue, wherein the first quantity of bandwidth is allocated to the first queue;

assigning a second quantity of bandwidth through the switching fabric to the first queue, wherein the second quantity of bandwidth is dynamically shared between the first queue and other queues;

passing cells associated with the first virtual connection through the switching fabric using the first quantity of bandwidth and at least a part of the second quantity of bandwidth; and

wherein the data cells comprise asynchronous transfer mode cells.

11. The method of claim 10, further comprising:

creating a first list of queues to prioritize access to dynamically shared bandwidth through the switching fabric; and

creating a second list of queues to prioritize access to dynamically shared bandwidth on the output link.

12. The method of claim 10, further comprising:

creating a plurality of prioritized lists of queues to prioritize access to dynamically shared bandwidth on the output link wherein each prioritized list corresponds to the type of traffic associated with the queues in the list;

creating a list of the plurality of prioritized lists of queues to further prioritize access to dynamically shared bandwidth on the output link.

13. The method of claim 4, further comprising:

controlling the flow of the first virtual connection from a switching fabric associated with the communications device to an output port of the communications device by utilizing the second queue associated with the first virtual connection.

14. The method of claim 5, further comprising:

controlling the flow of the first virtual connection from a switching fabric associated with the communications device to an output port of the communications device by utilizing the second queue associated with the first virtual connection; and

wherein the first queue is associated with a plurality of virtual connections and wherein the second queue is associated with a plurality of virtual connections.

15. The method of claim 14, wherein the data cells comprise asynchronous transfer mode cells.

16. The method of claim 14, further comprising:

assigning a first quantity of bandwidth through the switching fabric to the first queue, wherein the first quantity of bandwidth is allocated to the first queue;

assigning a second quantity of bandwidth through the switching fabric to the first queue, wherein the second quantity of bandwidth is dynamically shared between the first queue and other queues; and

passing cells associated with the first virtual connection through the switching fabric using the first quantity of bandwidth and at least a part of the second quantity of bandwidth.

17. The method of claim 14, further comprising:

assigning a third quantity of bandwidth on an output link of the communications device to the second queue, wherein the third quantity of bandwidth is allocated to the second queue;

assigning a fourth quantity of bandwidth on the output link to the second queue, wherein the fourth quantity of bandwidth is dynamically shared between the second queue and other queues; and

passing cells associated with the first virtual connection to the output link using the third quantity of bandwidth and at least a part of the fourth quantity of bandwidth.

18. The method of claim 17, further comprising:

assigning a first quantity of bandwidth through the switching fabric to the first queue, wherein the first quantity of bandwidth is allocated to the first queue;

assigning a second quantity of bandwidth through the switching fabric to the first queue, wherein the second quantity of bandwidth is dynamically shared between the first queue and other queues;

passing cells associated with the first virtual connection through the switching fabric using the first quantity of bandwidth and at least a part of the second quantity of bandwidth; and

wherein the data cells comprise asynchronous transfer mode cells.

19. The method of claim 18, further comprising:

creating a first list of queues to prioritize access to dynamically shared bandwidth through the switching fabric; and

creating a second list of queues to prioritize access to dynamically shared bandwidth on the output link.

20. The method of claim 18, further comprising:

creating a plurality of prioritized lists of queues to prioritize access to dynamically shared bandwidth on the output link wherein each prioritized list corresponds to the type of traffic associated with the queues in the list;

creating a list of the plurality of prioritized lists of queues to further prioritize access to dynamically shared bandwidth on the output link.

21. The method of claim 1 wherein the first queue comprises an output queue.

22. A method for controlling network service parameters in a cell based communications network, comprising:

receiving a plurality of input signals on input ports of a communications device, at least one input signal comprising at least one virtual connection, wherein a first virtual connection is associated with a first input port, wherein the communications device is part of the cell based communications network and wherein each virtual connection comprises a series of data cells comprising a header portion and a data portion;

assigning a first quantity of bandwidth through a switching fabric of the communications device to the first virtual connection, wherein the first quantity of bandwidth is allocated to the first virtual connection;

assigning a second quantity of bandwidth through the switching fabric to the first virtual connection, wherein the second quantity of bandwidth is dynamically shared between the first virtual connection and other virtual connections; and

passing cells associated with the first virtual connection through the switching fabric using the first quantity of bandwidth and at least a part of the second quantity of bandwidth.

23. The method of claim 22, wherein the data cells comprise asynchronous transfer mode cells.

24. The method of claim 23, further comprising:

assigning a third quantity of bandwidth on an output link of the communications device to the first virtual connection, wherein the third quantity of bandwidth is allocated to the first virtual connection;

assigning a fourth quantity of bandwidth on the output link to the first virtual connection, wherein the fourth quantity of bandwidth is dynamically shared between the first virtual connection and other virtual connections; and

passing cells associated with the first virtual connection to the output link using the third quantity of bandwidth and at least a part of the fourth quantity of bandwidth.

25. The method of claim 22, further comprising:

assigning a third quantity of bandwidth on an output link of the communications device to the first virtual connection, wherein the third quantity of bandwidth is allocated to the first virtual connection;

assigning a fourth quantity of bandwidth on the output link to the first virtual connection, wherein the fourth quantity of bandwidth is dynamically shared between the first virtual connection and other virtual connections; and

passing cells associated with the first virtual connection to the output link using the third quantity of bandwidth and at least a part of the fourth quantity of bandwidth.

26. An input port for a communications device, comprising:

a demultiplexer operable to receive asynchronous transfer mode cells associated with a plurality of virtual connections from a plurality of input links;

a buffer pool operable to store asynchronous transfer mode cells associated with the plurality of virtual connections received by the demultiplexer, the buffer pool including shared buffer space and dedicated buffer space;

a plurality of queues operable to store locations within the buffer pool of asynchronous transfer mode cells associated with the plurality of virtual connections, each of the plurality of queues being uniquely associated with a separate one of the plurality of virtual connections, a particular one of the plurality of queues being assigned by the demultiplexer to dedicated buffer space only available to that particular one of the plurality of queues, the particular one of the plurality of queues being assigned by the demultiplexer to the shared buffer space available to the particular one of the plurality of queues and other ones of the plurality of queues; and

a multiplexer operable to control transfer of asynchronous transfer mode cells associated with the plurality of virtual connections from the buffer pool.

27. The input port of claim 26, wherein the multiplexer is operable to allocate bandwidth to each of the plurality of queues in order to control transfer of asynchronous transfer mode cells associated with the plurality of virtual connections from the buffer pool.

28. The input port of claim 27, wherein a dedicated portion of the bandwidth is uniquely allocated by the multiplexer to the particular one of the plurality of queues.

29. The input port of claim 28, wherein a shared portion of the bandwidth is allocated by the multiplexer to be shared by the particular one of the plurality of queues and other ones of the plurality of queues.

30. The input port of claim 29, wherein the multiplexer assigns an available shared portion of the bandwidth allocated to be shared by two or more of the plurality of queues to the particular one of the plurality of queues upon exceeding its dedicated portion of the bandwidth.

31. The input port of claim 30, wherein the multiplexer prioritizes the plurality of queues to determine access capability to the shared portion of the allocated bandwidth.

32. The input port of claim 26, wherein the buffer pool is assigned to a specific traffic type.

33. The input port of claim 32, wherein the demultiplexer assigns available buffer locations of the shared buffer space to the particular one of the plurality of queues upon exceeding its number of buffer locations in its dedicated buffer space.

34. The input port of claim 26, wherein the dedicated buffer space and the shared buffer space have a plurality of buffer locations, each buffer location operable to store a single asynchronous transfer mode cell.

35. The input port of claim 34, wherein the demultiplexer prioritizes the plurality of queues to determine access capability to the shared buffer space.

36. A communication device, comprising:

an input port operable to receive asynchronous transfer mode cells associated with a plurality of virtual connections, the input port including a plurality of input queues, a particular one of the plurality of input queues being uniquely associated with a particular one of the plurality of virtual connections, the input port including a buffer pool having a dedicated input buffer space and a shared input buffer space to store asynchronous transfer mode cells, the dedicated input buffer space being uniquely associated with the particular one of the plurality of input queues, the shared input buffer space being associated with the particular one of the plurality of input queues and other ones of the plurality of input queues;

a switching fabric operable to transfer asynchronous transfer mode cells from the input port;

an output port operable to receive asynchronous transfer mode cells from the switching fabric, the output port including a plurality of output queues, a particular one of the plurality of output queues being uniquely associated with the particular one of the plurality of virtual connections, the output port operable to transfer received asynchronous transfer mode cells onto an output link, the output port including a buffer pool having a dedicated output buffer space and a shared output buffer space to store asynchronous transfer mode cells, the dedicated output buffer space being uniquely associated with the particular one of the plurality of output queues, the shared output buffer space being associated with the particular one of the plurality of output queues and other ones of the plurality of output queues.

37. The communication device of claim 36, wherein the input and output ports assign available shared input and output buffer space to the particular one of the plurality of input and output queues upon exceeding their dedicated buffer space, respectively.

38. The communication device of claim 37, wherein the input and output ports prioritize the plurality of input and output queues, to determine access capability to the shared input and output buffers, respectively.

39. The communication device of claim 36, wherein the input port uniquely allocates a dedicated portion of switching fabric bandwidth to the particular one of the plurality of input queues and allocates a shared portion of switching fabric bandwidth available for use by the particular one of the plurality of input queues and other ones of the plurality of input queues.

40. The communication device of claim 39, wherein the output port uniquely allocates a dedicated portion of output link bandwidth to the particular one of the plurality of output queues and allocates a shared portion of output link bandwidth available for use by the particular one of the plurality of output queues and one or more others of the plurality of output queues.

* * * * *